arXiv:1407.6288v2 [stat.ML] 24 Jul 2014

Subspace Learning From Bits

Yuejie Chi
Electrical and Computer Engineering
The Ohio State University
Columbus, OH 43204
Email: chi.97@osu.edu

July 25, 2014

Abstract

This paper proposes a simple sensing and estimation framework to faithfully recover the principal subspace of high-dimensional datasets or data streams from a collection of one-bit measurements from distributed sensors based on comparing accumulated energy projections of their data samples of dimension n over pairs of randomly selected directions. By leveraging low-dimensional structures, the top eigenvectors of a properly designed surrogate matrix is shown to recover the principal subspace of rank r as soon as the number of bit measurements exceeds the order of $nr^2 \log n$, which can be much smaller than the ambient dimension of the covariance matrix. The sample complexity to obtain reliable comparison outcomes is also obtained. Furthermore, we develop a low-complexity online algorithm to track the principal subspace that allows new bit measurements arrive sequentially. Numerical examples are provided to validate the proposed approach.

1 Introduction

An important challenge in modern data-intensive applications is to extract useful information without storing and transmitting all the data that may be generated distributively at an unprecedented rate. Fortunately, many practical datasets exhibit low-dimensional structures, such that a significant proportion of their variance can be captured in the first few principal components; and that the subspace spanned by these principal components is the recovery object of interest rather than the datasets themselves. Our goal in this paper is thus to design a simple yet efficient sensing and estimation framework to faithfully retrieve the *principal subspace* of high-dimensional datasets or data streams with low computation, communication and storage complexity.

Consider, for example, a data stream which generates a zero-mean data sample $\boldsymbol{x}_t \in \mathbb{R}^n$ at each time t. The covariance matrix of the data samples $\boldsymbol{\Sigma} = \mathbb{E}[\boldsymbol{x}_t \boldsymbol{x}_t^T]$ is assumed to be (approximately) low rank, $\text{rank}(\boldsymbol{\Sigma}) = r$, where $r \ll n$. This assumption is widely applicable to data such as network traffic, wideband spectrum, images, and so on. Several main challenges arise when processing such high-dimensional data streams.

- First, each data sample $\boldsymbol{x}_t$ may not be fully observed, but only sketched once or twice via a linear inner product. This may arise either due to data-on-the-fly, or privacy concerns so that users do not reveal the complete samples. This requires the computation to be accomplished with only one pass or two passes of the data sample with a low computational complexity.

- Secondly, the data samples $\boldsymbol{x}_t$ may arrive at distributed sensors (or machines) to first obtain a few measurements locally, and then these measurements are communicated to a central *fusion center* to compute the principal subspace of the covariance matrix without transmitting the original data samples. The distributed sensors usually are power-hungry and resource-limited, therefore it is highly desirable to minimize the computational and storage cost at the sensor side, as well as the communication overheads to the fusion center.

- Thirdly, as new data samples arrive and/or new sensors enter, it is interesting to *track* the changes of the principal subspace at the fusion center without storing all history data.

This paper proposes a simple framework to deal with the above challenges with provable performance guarantees. In particular, we focus on the scenario where each sensor, possibly distributed, may only access a subset or substream of the whole data, and transmits their measurements to the fusion center for subspace estimation. It is of interest to reduce the amount of measurements required to communicate to the fusion center while keeping the number of sensors as small as possible to allow faithful recovery of the principal subspace.

At the sensing stage, each sensor computes the accumulated energy projections of the data samples onto *two* randomly selected directions with i.i.d. Gaussian entries, called *sketching vectors*. It then transmits a *one-bit* comparison outcome to the fusion center indicating the direction onto which it has a higher energy projection. This corresponds to comparing the energy projection of the *sample covariance matrix* onto two randomly selected rank-one subspaces. A key observation is that as long as the number of samples is not too small, the comparison outcome will be exactly the same as if it is performed on the *population covariance matrix*. By only transmitting the comparison outcome rather than the actual energy measurements, the communication overhead is minimized to a single bit and the measurements are more robust to communication channel errors. Moreover, as will be shown, the energy projections can be computed extremely simple without storing history data samples, and are always nonnegative, making them suitable for wideband and optical applications.

At the fusion center, the principal subspace of the covariance matrix is estimated as the top eigenvectors of a carefully designed surrogate matrix using the collected one-bit measurements. The top eigenvectors can be easily computed via a truncated Eigen Value Decomposition (EVD) or power iteration with low complexity. The batch computation requires the fusion center to assume knowledge of all sketching vectors, which may be infeasible. To ease this requirement, we increase the communication overhead by requiring the sensors to transmit both the one-bit measurement and their sketching vectors (or their seed) to the fusion center. We then developed a memory-efficient algorithm to online update the estimate of the principal subspace assuming the measurements arrive sequentially, which can be implemented without storing all the sketching vectors and have a memory requirement on the order of the size of the principal subspace.

Two quantities are of importance for the proposed framework to succeed: the total number of one-bit measurements (or sensors), m; and the number of samples seen by each sensor, T. We show that, assume all the bit measurements are exact and the rank r is fixed, the estimate of the principal subspace is provably accurate as long as m is on the order of $nr^2 \log n$, when the projection directions are composed of i.i.d. Gaussian entries. This demonstrates a significant gain of exploiting low-dimensional structures. We then show that for a fixed $\boldsymbol{\Sigma}$, as long as T is on the order of $(\mathrm{Tr}(\boldsymbol{\Sigma})/\|\boldsymbol{\Sigma}\|_{\mathrm{F}} \log^2 m)$ samples are seen per sensor (regardless whether they observe the same set of data samples or not), the bit measurements are accurate with high probability.

1.1 Related Work

Estimating a low-dimensional data stream from its sparse observations has attracted interest in recent years, but most existing work has been assuming a sampling rate that allows recovering of the data stream [1, 2]. The proposed framework is motivated by our recent work in [3, 4, 5], where a quadratic sampling scheme is designed for low-rank covariance estimation. It is shown in [3] that a number of $\Theta(nr)$ quadratic (energy) measurements suffices to exactly recover a rank-r covariance matrix via nuclear norm minimization, assume the measurement vectors are composed of i.i.d. sub-Gaussian entries. However, communicating these energy measurements with high precision may cause unwanted overhead in practice. Though tractable, the convex optimization algorithm in [3] is still computationally expensive when the data dimensionality is high and requires estimation of the noise level for implementation.

A related line of research is on one-bit compressed sensing [6, 7, 8, 9], where the authors aim to recover the signal with low-dimensional structures from the signs of random linear measurements up to a scaling factor. In contrast, our measurements are signs of random *quadratic* measurements. In this paper, we also briefly consider a convex optimization algorithm for the fusion center which is revised straightforwardly from the algorithm proposed by Plan and Vershynin [9] for one-bit compressed sensing. However, the performance guarantee therein [9] cannot be directly applied due to the nonlinear measurements. Nonetheless, we demonstrate through numerical examples that the revised convex optimization program has a similar performance with the much simpler truncated EVD, in terms of recovery accuracy and robustness to flipping errors. Therefore, we advocate the simpler approach using truncated EVD for the fusion center.

Distributed wideband spectrum sensing for cognitive radios [10] is an appealing and motivating application for the proposed framework. It is recently proposed to estimate the power spectral density of wideband signals via least-squares estimation from sub-Nyquist samples [11]. The frugal sensing framework [12] considered the same estimation problem using one-bit measurements based on comparing the average signal power within a band of interest against a pre-determined or adaptively-set threshold [13]. Their algorithm is based on linear programming and may explore parametric representations of the power spectral density. Our work is different from [12] in several aspects. Instead of comparing the average signal power against a threshold which may introduce additional parameters, we compare the average signal power between two different bands of interest. Our algorithm explores the low-rank property of the power spectral density rather than its parametric representation.

Finally, the paper [14] studied one-bit phase retrieval, an extension of the one-bit compressed sensing with phaseless measurements. Despite different motivations and applications, our algorithm subsumes the scenario in [14] as a special case when the covariance matrix is assumed rank-one.

1.2 Paper Organization

The rest of this paper is organized as follows. Section 2 describes the sensing framework and formulates the subspace estimation problem using one-bit measurements. Section 3 presents the performance guarantees of the proposed scheme. Section 4 briefly discussed an alternative recovery algorithm based on convex optimization that is amenable to other covariance structures. Section 5 presents an online algorithm to track the low-rank principal subspace with sequential bit measurements. Numerical examples are given in Section 6. Finally, we conclude in Section 7 and outline some future directions.

2 Subspace Estimation From One-Bit Energy Comparisons

In this section we describe the distributed sensing and central estimation framework for subspace estimation utilizing one-bit energy comparisons from each sensor, summarized in Algorithm 1.

2.1 Signal Model

Let $\boldsymbol{x}_t \in \mathbb{R}^n$ be a sample in the dataset or data stream with $\mathbb{E}[\boldsymbol{x}_t] = 0$ and its covariance matrix $\boldsymbol{\Sigma} = \mathbb{E}[\boldsymbol{s}_t \boldsymbol{s}_t^T] = \sum_{k=1}^r \lambda_k \boldsymbol{u}_k \boldsymbol{u}_k^T = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^T \in \mathbb{R}^{n\times n}$, with $\boldsymbol{U} = [\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_r]$. The covariance matrix $\boldsymbol{\Sigma}$ is a low-rank positive semidefinite (PSD) matrix with $\operatorname{rank}(\boldsymbol{\Sigma}) = r \ll n$ and $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_r$. We further define $\boldsymbol{v}_k$, $k = 1, \ldots, n-r$ as the basis vectors spanning the complement of $\boldsymbol{u}_k$, $k = 1, \ldots, r$. Our goal is to recover $\boldsymbol{U} \in \mathbb{R}^{n\times r}$.

2.2 One-Bit Sensing Scheme

Consider a collection of m sensors that are deployed distributively to measure the data stream[1]. Each sensor can either access the complete data stream or a substream. At the ith sensor, define a pair of sketching vectors $\boldsymbol{a}_i \in \mathbb{R}^n$ and $\boldsymbol{b}_i \in \mathbb{R}^n$, $1 \leq i \leq m$, where their entries are i.i.d. standard Gaussian. Without loss of generality, we assume the data samples accessed by the ith sensor are indexed by $\{\ell_t^i\}_{t=1}^T$. At the tth sample time, the ith sensor takes the quadratic measurements:

$$u_{i,t} = |\langle \boldsymbol{a}_i, \boldsymbol{x}_{\ell_t^i} \rangle|^2, \quad v_{i,t} = |\langle \boldsymbol{b}_i, \boldsymbol{x}_{\ell_t^i} \rangle|^2,$$

which are phaseless and more reliable in high frequency applications. These quadratic measurements are then averaged over T samples to obtain

$$U_{i,T} = \frac{1}{T}\sum_{t=1}^T u_{i,t} = \boldsymbol{a}_i^T \boldsymbol{\Sigma}_{i,T} \boldsymbol{a}_i, \; V_{i,T} = \frac{1}{T}\sum_{t=1}^T v_{i,t} = \boldsymbol{b}_i^T \boldsymbol{\Sigma}_{i,T} \boldsymbol{b}_i,$$

where

$$\boldsymbol{\Sigma}_{i,T} = \frac{1}{T}\sum_{t=1}^T \boldsymbol{x}_{\ell_t^i} \boldsymbol{x}_{\ell_t^i}^T$$

is the sample covariance matrix seen by the ith sensor. It is clear that $U_{i,T} = \frac{T-1}{T} U_{i,T-1} + \frac{1}{T} u_{i,T}$, and similarly $V_{i,T}$, can be updated recursively without storing all the history data. At the end of sample time T, the ith sensor compares the average energy $U_{i,T}$ and $V_{i,T}$, and transmit to the fusion center a single bit $y_{i,T}$ indicating the outcome:

$$y_{i,T} = \begin{cases} 1, & \text{if } U_{i,T} > V_{i,T} \\ -1, & \text{otherwise} \end{cases}. \tag{1}$$

Depending on whether the fusion center has prior knowledge of the sketching vectors, the sensors may also need to transmit $\boldsymbol{a}_i$ and $\boldsymbol{b}_i$. The communication overhead is minimized when only the bit indicating the comparison outcome is transmitted.

Since the sample covariance matrix $\boldsymbol{\Sigma}_T$ converges to the true covariance matrix as T approaches infinity, the bit measurement at the ith sensor also approaches to the following as soon as T is not too large (as shown later)

$$y_i = \begin{cases} 1, & \text{if } a_i^T \boldsymbol{\Sigma} a_i > b_i^T \boldsymbol{\Sigma} b_i \\ -1, & \text{otherwise} \end{cases}. \tag{2}$$

[1]The approach for a batch data set is straightforward.

Algorithm 1 Subspace Estimation From One-Bit Energy Comparisons

1: **for** each sensor $i = 1, \ldots, m$ **do**

2: Randomly choose two sketch vectors $\boldsymbol{a}_i \in \mathbb{R}^n$ and $\boldsymbol{b}_i \in \mathbb{R}^n$ with i.i.d. Gaussian entries;

3: Sketch an arbitrary subset (or substream) indexed by $\{\ell_t^i\}_{t=1}^T$ with two energy measurements $|\langle \boldsymbol{a}_i, \boldsymbol{x}_{\ell_t^i}\rangle|^2$ and $|\langle \boldsymbol{b}_i, \boldsymbol{x}_{\ell_t^i}\rangle|^2$, and transmit a binary bit to the fusion center:

$$y_{i,T} = \operatorname{sign}\left(\frac{1}{T}\sum_{t=1}^{T} |\langle \boldsymbol{a}_i, \boldsymbol{x}_{\ell_t^i}\rangle|^2 - \frac{1}{T}\sum_{t=1}^{T} |\langle \boldsymbol{b}_i, \boldsymbol{x}_{\ell_t^i}\rangle|^2\right)$$

which can be implemented in a manner without storing all history data.

4: **end for**

5: **Estimation:** The fusion center recovers the principal components $\hat{\boldsymbol{U}} \in \mathbb{R}^{n\times r}$ by computing the top r eigenvectors of the surrogate matrix $\boldsymbol{J}_T = \frac{1}{m}\sum_{i=1}^m y_{i,T}\left(\boldsymbol{a}_i\boldsymbol{a}_i^T - \boldsymbol{b}_i\boldsymbol{b}_i^T\right)$.

Let $\boldsymbol{W}_i = \boldsymbol{a}_i\boldsymbol{a}_i^T - \boldsymbol{b}_i\boldsymbol{b}_i^T$, we can then compactly represent (1) and (2) as

$$y_{i,T} = \operatorname{sign}\left(\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}_{i,T}\rangle\right) \quad \text{and} \quad y_i = \operatorname{sign}\left(\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle\right), \tag{3}$$

where $\operatorname{sign}(\cdot)$ is the sign function. The above operation can be regarded as comparing the energy projections of the covariance matrix onto two randomly selected directions.

Apparently not all information about $\boldsymbol{\Sigma}$ can be recovered from the bit measurements $\{y_i\}_{i=1}^m$. For example, let $\boldsymbol{\Sigma} = \lambda\boldsymbol{C}\boldsymbol{C}^T$ be a rank-r matrix with $\boldsymbol{C} \in \mathbb{R}^{n\times r}$ and $\lambda = \|\boldsymbol{\Sigma}\|_{\mathrm{F}}$. Apparently,

$$\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i = \lambda\|\boldsymbol{C}^T\boldsymbol{a}_i\|_2^2 = \lambda\|\boldsymbol{Q}\boldsymbol{C}^T\boldsymbol{a}_i\|_2^2,$$

for $\boldsymbol{Q} \in \mathbb{O}^r$, hence only the subspace spanned by $\boldsymbol{C}$ can be recovered up to orthonormal transformations. Moreover, the value of λ cannot be recovered either. Therefore, for the rest of the paper we focus on recovering the principal subspace of $\boldsymbol{\Sigma}$ from the one-bit measurements $\{y_{i,T}\}_{i=1}^m$ given in (1).

2.3 Principal Subspace Estimation via Truncated EVD

We propose an extremely simple and low-complexity estimator whose complexity amounts to computing a few top eigenvectors. For will first consider subspace estimation from $\{y_i\}_{i=1}^m$ as if the bit measurements To motivate the algorithm, consider the simple case when $\boldsymbol{\Sigma} = \boldsymbol{\nu}\boldsymbol{\nu}^T$ is a rank-one matrix with $\|\boldsymbol{\nu}\|_2 = 1$. A natural way to recover $\boldsymbol{\nu}$ is from the following optimization:

$$\max_{\boldsymbol{\nu}} \frac{1}{m}\sum_{i=1}^m y_{i,T}\left\langle \boldsymbol{W}_i, \boldsymbol{\nu}\boldsymbol{\nu}^T\right\rangle = \left\langle \frac{1}{m}\sum_{i=1}^m y_{i,T}\boldsymbol{W}_i, \boldsymbol{\nu}\boldsymbol{\nu}^T\right\rangle \qquad \text{s.t.} \quad \|\boldsymbol{\nu}\|_2 = 1, \tag{4}$$

which aims to find the direction $\boldsymbol{\nu}$ that agrees with the measured signs as much as possible. Clearly, the solution to (4) is the top eigenvector of the surrogate matrix:

$$\boldsymbol{J}_T = \frac{1}{m}\sum_{i=1}^m y_{i,T}\boldsymbol{W}_i. \tag{5}$$

More generally, we recover the principal subspace of a rank-r covariance matrix $\boldsymbol{\Sigma}$ as the top-r eigenvectors of the above matrix (5), which we denoted as the truncated EVD approach.

3 Performance Analysis

When the bit measurements $y_{i,T}$'s all agree with y_i's given in (2), we have $\boldsymbol{J}_T$ in agreement with the following

$$\boldsymbol{J} = \frac{1}{m}\sum_{i=1}^{m} y_i \boldsymbol{W}_i = \frac{1}{m}\sum_{i=1}^{m} \mathrm{sign}(\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle)\boldsymbol{W}_i.$$

In the sequel we establish the performance guarantee of the proposed scheme in two steps. First, we show the rank-r subspace of $\boldsymbol{\Sigma}$ can be accurately estimated using $\boldsymbol{J}$ as soon as m is slightly above the information-theoretical limit. Next, we give sample complexity on T establishing the condition that $y_{i,T}$'s all agree with y_i's.

3.1 Number of Required Bit Measurements

Define $\bar{\boldsymbol{J}} = \mathbb{E}[\boldsymbol{J}]$. We first characterize the eigenspace of $\bar{\boldsymbol{J}}$ and then show that $\boldsymbol{J}$ is concentrated around $\bar{\boldsymbol{J}}$ for sufficiently large m. First, we have the following lemma whose proof is provided in Appendix B.

Lemma 1. *The principal subspace of $\bar{\boldsymbol{J}}$ is the same as that of $\boldsymbol{\Sigma}$ with rank$(\bar{\boldsymbol{J}}) = r$.*

$$1 \geq \boldsymbol{u}_k^T \bar{\boldsymbol{J}} \boldsymbol{u}_k \geq \max\left\{ \left(\frac{1}{1+\kappa(\boldsymbol{\Sigma})}\right)^{r-1}, \frac{1}{9r} e^{-\kappa(\boldsymbol{\Sigma})} \right\} := \alpha, \quad k = 1, \ldots, r; \tag{6}$$

$$\boldsymbol{v}_k^T \bar{\boldsymbol{J}} \boldsymbol{v}_k = 0, \quad k = 1, \ldots, n-r, \tag{7}$$

where $\kappa(\boldsymbol{\Sigma}) = \lambda_1/\lambda_k$ is the conditioning number of $\boldsymbol{\Sigma}$. When $r = 1$, the right-hand side of (6) equals one.

Following Lemma 1, the spectral gap between the rth eigenvalue and the $r+1$th eigenvalue (which is zero) of $\bar{\boldsymbol{J}}$ is at least

$$\max\left\{ \left(\frac{1}{1+\kappa(\boldsymbol{\Sigma})}\right)^{r-1}, \frac{1}{9r} e^{-\kappa(\boldsymbol{\Sigma})} \right\},$$

where the first term (exponential lower bound) is tighter when r is small while the second term (polynomial lower bound) is more accurate when r is large. Indeed $\boldsymbol{u}_k^T \bar{\boldsymbol{J}} \boldsymbol{u}_k$ may be computed exactly and we compare it with the derived lower bounds in Fig. 1 by assuming all $\lambda_k = 1$. Indeed the polynomial bound is rather accurate except the leading constant.

Next, we need to show that for sufficiently large m, the matrix $\boldsymbol{J}$ is close to its expectation $\bar{\boldsymbol{J}}$ using the matrix Bernstein's inequality in Lemma 4. We have the following lemma whose proof can be found in Appendix C.

Lemma 2. *Let $\boldsymbol{\Sigma}$ be a fixed matrix. Let $0 < \delta < 1$. Then with probability at least $1-\delta$, we have that*

$$\|\boldsymbol{J} - \bar{\boldsymbol{J}}\| \leq \tau,$$

where $\tau = \sqrt{\frac{c_1 n}{m} \log\left(\frac{2n}{\delta}\right)}$, where c_1 is an absolute constant.

Our main theorem then immediately follows by an improvement of the Davis-Kahan sin-Theta theorem [15] in Lemma 5. Denote the principal subspace of $\boldsymbol{J}$ as $\hat{\boldsymbol{U}}$, we have the following theorem.

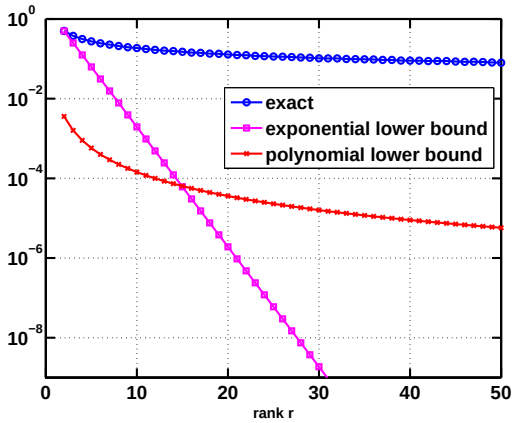

Figure 1: The derived lower bounds in Lemma 1 compared with its exact value when all $\lambda_k = 1$ for different ranks.

Theorem 1. *Let $\boldsymbol{\Sigma}$ be a fixed matrix, and $0 < \delta < 1$. With probability at least $1 - \delta$, there exists a $r \times r$ orthogonal matrix $\boldsymbol{Q}$ such that*

$$\begin{aligned}\|\hat{\boldsymbol{U}} - \boldsymbol{U}\boldsymbol{Q}\|_{\mathrm{F}} &\leq \alpha^{-1}\sqrt{\frac{c_1 n}{m}\log\left(\frac{2n}{\delta}\right)} \\ &\leq \min\left\{(1+\kappa(\boldsymbol{\Sigma}))^{r-1}, 9e^{\kappa(\boldsymbol{\Sigma})}r\right\}\sqrt{\frac{c_1 n}{m}\log\left(\frac{2n}{\delta}\right)},\end{aligned}$$

where c_1 is an absolute constant, and $\kappa(\boldsymbol{\Sigma})$ is its conditioning number.

Remark: As shown earlier, for not too small r and small $\kappa(\boldsymbol{\Sigma})$, the polynomial bound in α dominates and we have that as soon as the number of bit measurements exceeds the order of $\Theta(nr^2 \log n)$, it is sufficient to recover to recover the principal subspace $\boldsymbol{U}$ with high accuracy. Since there are at least nr degrees of freedom to describe the subspace $\boldsymbol{U}$, our bound is near-optimal up to a polynomial factor with respect to r and a logarithmic factor with respect to n. It is worth emphasizing that a highlight of our result indicates that a number of *bit measurements* that is much smaller than the ambient dimension of the covariance matrix could still yield reliable subspace estimate.

3.2 Sample Complexity to Guarantee Faithful Bit Measurements

In practice, instead of directly measuring the energy projection of the covariance matrix as in (2), we measure the energy projection of the sample covariance matrix in (1). Therefore it is necessary to derive the sample complexity required to obtain faithful bit measurements through (1). Note that the measured energy difference before taking its sign is

$$z_i = \langle \boldsymbol{W}_i, \boldsymbol{\Sigma}_T \rangle = \frac{1}{T}\sum_{t=1}^{T} \langle \boldsymbol{W}_i, \boldsymbol{x}_t \boldsymbol{x}_t^T \rangle \tag{8}$$

where $\boldsymbol{\Sigma}_{i,T} = \boldsymbol{\Sigma}_T$ and $\ell_t^i = t$ for simplicity. In this subsection, we derive how large T needs to be in order to guarantee accurate bit measurements. We have the following proposition whose proof is in Appendix D.

Proposition 1. *Let $0 < \delta \leq 1$. Assume $\boldsymbol{x}_t$ are i.i.d. Gaussian satisfying $\boldsymbol{x}_t \sim \mathcal{N}(0, \boldsymbol{\Sigma})$. Then*

$$\mathbb{P}\left[z_i \neq sign(\langle \boldsymbol{W}_i, \boldsymbol{\Sigma} \rangle)\right] \leq \delta$$

as soon as $T > c\frac{\mathrm{Tr}(\boldsymbol{\Sigma})}{\|\boldsymbol{\Sigma}\|_{\mathrm{F}}} \log^2(1/\delta)$ for some sufficiently large constant c.

In order to guarantee that all m bits are accurately estimated, we need to further apply a union bound on Proposition 1, which yields the following theorem.

Theorem 2. *Let $\boldsymbol{x}_t$ be zero-mean with $\boldsymbol{x}_t \sim \mathcal{N}(0, \boldsymbol{\Sigma})$. Let $0 < \delta \leq 1$. With probability at least $1-\delta$, all bit measurements are exact, given that the number of samples observed by each sensor satisfies*

$$T > c\frac{\mathrm{Tr}(\boldsymbol{\Sigma})}{\|\boldsymbol{\Sigma}\|_{\mathrm{F}}} \log^2\left(\frac{m}{\delta}\right)$$

for some sufficiently large constant c.

Remark: Recall that $\mathrm{Tr}(\boldsymbol{\Sigma}) \leq \sqrt{r}\|\boldsymbol{\Sigma}\|_{\mathrm{F}}$, then as soon as T is on the order of $\sqrt{r}\log^2 m$ all bit measurements are accurate with high probability. Combining with Theorem 1, if the number of samples grows exceeds the order of $\min\{r^{\frac{1}{2}} \log^2(nr^2 \log n), r^{\frac{3}{2}} \log^2(n \log n)\}$ which only depends on the ambient dimension logarithmically assuming all sensors are measuring the same set of samples.

4 An Alternative Convex Program

The performance guarantees of the above truncated EVD approach hold when the covariance matrix $\boldsymbol{\Sigma}$ is exactly low-rank. In practice, $\boldsymbol{\Sigma}$ may be approximately low-rank, which can be captured by the following constraint [9]:

$$\mathcal{K} = \{\boldsymbol{\Sigma} : \|\boldsymbol{\Sigma}\|_{\mathrm{F}} \leq 1,\ \|\boldsymbol{\Sigma}\|_* \leq \sqrt{r}\}.$$

Plan and Vershynin [9] studied the general problem of estimating a signal $\boldsymbol{\Sigma} \in \mathcal{K}$ when $\boldsymbol{W}_i$'s are composed of i.i.d. Gaussian entries. Specializing their algorithm to low-rank PSD matrices, we can consider the following convex optimization algorithm to recover $\boldsymbol{\Sigma}$:

$$\hat{\boldsymbol{\Sigma}} = \arg\max_{\boldsymbol{\Sigma} \succeq 0} \sum_{i=1}^{m} y_i \langle \boldsymbol{W}_i, \boldsymbol{\Sigma} \rangle, \quad \text{s.t.} \quad \boldsymbol{\Sigma} \in \mathcal{K}. \tag{9}$$

Unfortunately the performance guarantees in [9] do not apply anymore due to the non-independence of the entries of $\boldsymbol{W}_i$ in our case. Nonetheless, we evaluate the performance of (9) numerically in the Section 6, and show its performance is comparable to that of the truncated EVD approach, while incurring a much higher computational cost.

It is worth mentioning that the convex optimization algorithm are amenable to other types of low-dimensional covariance structures. In some applications, the covariance matrix may be modeled as approximately sparse, captured by the constraint:

$$\mathcal{K} = \{\boldsymbol{\Sigma} : \|\boldsymbol{\Sigma}\|_{\mathrm{F}} \leq 1,\ \|\boldsymbol{\Sigma}\|_1 \leq \sqrt{s}\},$$

where s is the sparsity level of $\boldsymbol{\Sigma}$ and $\|\cdot\|_1$ is the element-wise ℓ_1 norm. The algorithm (9) can be straightforwardly revised by modifying the constraint set.

5 Subspace Tracking with Online Bit Measurements

In this section, we develop an online subspace estimation algorithm for the fusion center to update the principal subspace estimate as new bit measurements arrive. This is particularly useful when the fusion center is also *memory limited* so that it is incapable to store all sketching vectors up to a memory space of $\Theta(n^2)$. The proposed online algorithm has a memory space of $\Theta(nr)$, that is the size of the principal subspace, and does not require storing all the sketching vectors. In the online setting, the sensors transmit both the bit measurement and the sketching vectors, while the fusion center assumes no prior knowledge about the sketching vectors.

We now describe an online algorithm to estimate the principal subspace. Essentially we need to update the principal subspace of $\boldsymbol{J}_m$ from that of $\boldsymbol{J}_{m-1}$ given a rank-two update as follows:

$$\boldsymbol{J}_m = \frac{m-1}{m}\boldsymbol{J}_{m-1} + \frac{y_m}{m}\left(\boldsymbol{a}_m\boldsymbol{a}_m^T - \boldsymbol{b}_m\boldsymbol{b}_m^T\right) \tag{10}$$

We can rewrite (10) using more general notations as

$$\boldsymbol{J}_m = \eta_m\boldsymbol{J}_{m-1} + \boldsymbol{K}_m\boldsymbol{\Lambda}_m\boldsymbol{K}_m^T, \tag{11}$$

where (10) can be obtained from (11) by letting $\eta_m = \frac{m-1}{m}$, $\boldsymbol{K}_m = [\boldsymbol{a}_m, \boldsymbol{b}_m] \in \mathbb{R}^{n\times 2}$, and $\boldsymbol{\Lambda}_m = \mathrm{diag}([y_m/m, -y_m/m])$. Note that it might be of interest to incorporate an additional discounting factor on η_m to emphasize the current measurement, by letting η_m to take a smaller value, as done in [16], but we do not find it particularly helpful in our problem.

Assume the EVD of $\boldsymbol{J}_{m-1}$ can be written as $\boldsymbol{J}_{m-1} = \boldsymbol{U}_{m-1}\boldsymbol{\Pi}_{m-1}\boldsymbol{U}_{m-1}^T$ where $\boldsymbol{U}_{m-1} \in \mathbb{R}^{n\times r}$ is orthonormal and $\boldsymbol{\Pi}_{m-1} \in \mathbb{R}^{r\times r}$ is diagonal. The goal is to find the best rank-r approximation of $\boldsymbol{J}_m$ by updating $\boldsymbol{U}_{m-1}$ and $\boldsymbol{\Pi}_{m-1}$.

We develop a fast rank-two update of the EVD of a symmetric matrix by introducing necessary modifications of the incremental SVD approach in [16, 17]. A key difference from [16, 17] is that we do not allow the size of the principal subspace to grow, which is fixed as r. In the update we first compute an expanded principal subspace of rank $(r+2)$ and then only keep its r largest principal components.

Let $\boldsymbol{R}_m = (\boldsymbol{I} - \boldsymbol{U}_{m-1}\boldsymbol{U}_{m-1}^T)\boldsymbol{K}_m$ and $\boldsymbol{P}_m = \mathrm{orth}(\boldsymbol{R}_m)$ be the orthonormal columns spanning $\boldsymbol{R}_m$. We can write $\boldsymbol{J}_m$ as

$$\begin{aligned}\boldsymbol{J}_m &= \begin{bmatrix}\boldsymbol{U}_{m-1} & \boldsymbol{P}_m\end{bmatrix}\left(\begin{bmatrix}\eta_m\boldsymbol{\Pi}_{m-1} & 0\\ 0 & 0\end{bmatrix} + \begin{bmatrix}\boldsymbol{U}_{m-1}^T\boldsymbol{K}_m\\ \boldsymbol{P}_m^T\boldsymbol{R}_m\end{bmatrix}\boldsymbol{\Lambda}_m\begin{bmatrix}\boldsymbol{U}_{m-1}^T\boldsymbol{K}_m\\ \boldsymbol{P}_m^T\boldsymbol{R}_m\end{bmatrix}^T\right)\begin{bmatrix}\boldsymbol{U}_{m-1}^T\\ \boldsymbol{P}_m^T\end{bmatrix}\\ &:= \begin{bmatrix}\boldsymbol{U}_{m-1} & \boldsymbol{P}_m\end{bmatrix}\boldsymbol{\Gamma}_m\begin{bmatrix}\boldsymbol{U}_{m-1}^T\\ \boldsymbol{P}_m^T\end{bmatrix},\end{aligned}$$

where $\boldsymbol{\Gamma}_m$ is a small $(r+2)\times(r+2)$ matrix whose EVD can be computed easily and yields

$$\boldsymbol{\Gamma}_m = \boldsymbol{U}'_m\boldsymbol{\Pi}'_m\boldsymbol{U}'_m.$$

Set $\boldsymbol{\Pi}_m$ be the top $r\times r$ sub-matrix of $\boldsymbol{\Pi}'_m$ assuming the eigenvalues are given in an absolute descending order, the principal subspace of $\boldsymbol{J}_m$ can be updated correspondingly as

$$\boldsymbol{U}_m := \begin{bmatrix}\boldsymbol{U}_{m-1} & \boldsymbol{P}_m\end{bmatrix}\boldsymbol{U}'_m\boldsymbol{I}_r$$

where $\boldsymbol{I}_r$ is the first r columns of the $(r+2)\times(r+2)$ identity matrix.

6 Numerical Experiments

In the numerical experiments, we first examine the performance of the truncated EVD algorithm in a batch setting in terms of reconstruction accuracy and robustness to flipping errors, and compare it against the convex optimization algorithm in (9). We then examine the performance of the online subspace estimation algorithm in Section 5 and apply it to the problem of line spectrum estimation.

6.1 Truncated-EVD Based Recovery

We generate the covariance matrix as $\boldsymbol{\Sigma} = \boldsymbol{X}\boldsymbol{X}^T$, where $\boldsymbol{X} \in \mathbb{R}^{n\times r}$ is composed of standard Gaussian entries. The sketching vectors $\boldsymbol{a}_i$'s and $\boldsymbol{b}_i$'s are also generated with standard Gaussian entries. After the bit measurements are collected, we run the truncated EVD algorithm and the convex optimization algorithm (9) assuming the rank r of principal subspace is known perfectly. The algorithm (9) is performed using the MOSEK toolbox available in CVX and obtained an estimate $\hat{\boldsymbol{\Sigma}}$, from which we extract its top r eigenvectors $\hat{\boldsymbol{X}}$. The truncated EVD algorithm directly obtains an estimate of $\hat{\boldsymbol{X}}$. The error metric is calculated as the Normalized Mean Squared Error (NMSE) given as $\|(\boldsymbol{I}-\hat{\boldsymbol{X}}\hat{\boldsymbol{X}}^T)\boldsymbol{X}\|_{\mathrm{F}}^2/\|\boldsymbol{X}\|_{\mathrm{F}}^2$. Fig. 2 shows the average NMSE with respect to the number of bit measurements for different dimensions $n = 40, 100, 200$ when $r = 3$ averaged over 10 Monte Carlo runs. Given the high complexity of the convex algorithm (9), we only perform it when $n = 40$. We can see that its performance is comparable to that of the truncated EVD. Fig. 3 further examines the performance of the truncated EVD algorithm for different ranks. For the same number of bit measurements, the NMSE grows gracefully as the rank increases.

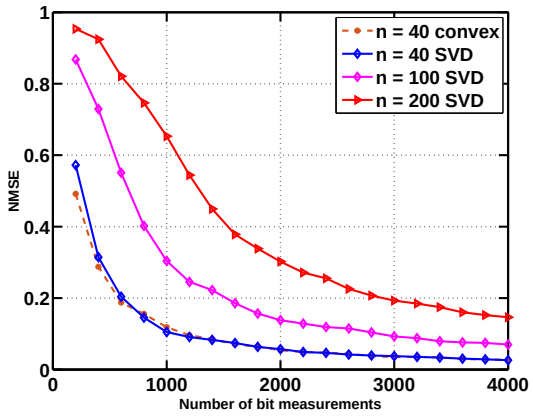

Figure 2: NMSE with respect to the number of bit measurements for estimating the principal subspace of the covariance matrix with $n = 40, 100, 200$ when $r = 3$.

The truncated EVD algorithm also exhibits a reasonable robustness against flipping errors. We assume that each bit measurement may be flipped independently with a flipping probability $0 \le \epsilon \le 1/2$. Fig. 4 shows the reconstructed NMSE with respect to the flipping probability for different number of bit measurements m when $n = 100$ and $r = 3$.

6.2 Online Subspace Tracking and Application to Line Spectrum Estimation

We now examine the performance of the online subspace estimation algorithm proposed in Section 5. Let $n = 40$ and $r = 3$. Fig. 5 shows the NMSE of principal subspace estimation with respect to the

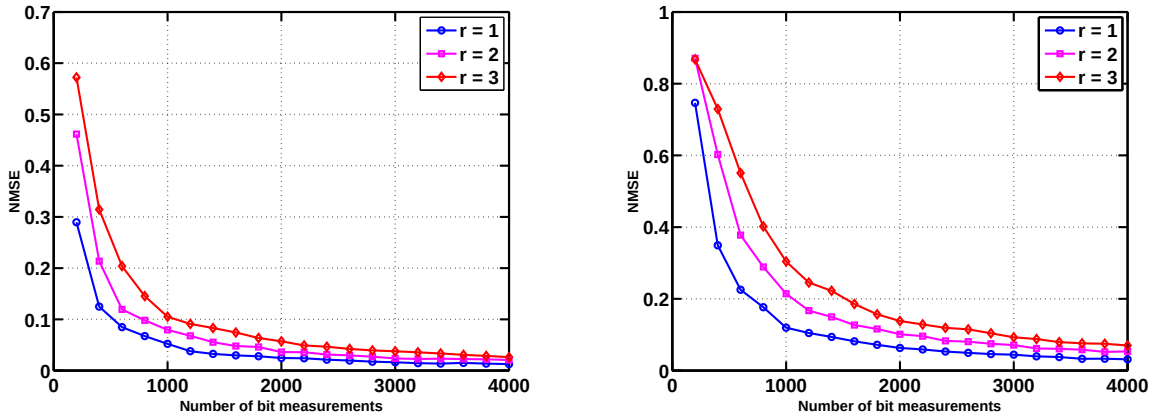

(a) $n = 40$ (b) $n = 100$

Figure 3: NMSE of principal subspace estimation via the truncated EVD algorithm with respect to the number of bit measurements when (a) $n = 40$ and (b) $n = 100$ with different ranks.

number of bit measurements, where the result is averaged over 10 Monte Carlo runs. Compared with Fig. 3 (a), the estimation accuracy is comparable to that in a batch setting.

In the sequel, we consider the problem of line spectrum estimation from its power spectral density, where we assume the covariance matrix $\boldsymbol{\Sigma}$ is a low-rank Toeplitz PSD matrix. Let $n = 40$ and $r = 3$. By the Vandermonde decomposition lemma [18], we write $\boldsymbol{\Sigma} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^H$, where $\boldsymbol{V} \in \mathbb{C}^{n\times r}$ is a Vandermonde matrix composed of the frequencies in a set $\mathcal{F}$, and $\boldsymbol{\Lambda} = \mathrm{diag}(\sigma^2)$ is a diagonal matrix describing the variance of each frequency.

At each new measurement, we first use the online subspace estimation algorithm proposed in Section 5 to estimate a principal subspace of rank $r_{\text{est}} = 5$, then apply ESPRIT [19] to the reconstructed principal subspace to recover the set of frequencies. Fig. 6 shows the estimation results for various parameter settings, where the estimates of frequency locations are plotted vertically at each new bit measurement. The color indicates the amplitudes of the estimated frequencies. Fig. 6 (a) and (b) have the same frequency profile $\mathcal{F}$, with two close located frequencies separated by the Rayleigh limit. When all the frequencies have strong amplitudes, the frequencies can be accurately estimated as depicted in (a); when one of the close frequency is relatively weak, the algorithm requires more measurements to pick up the weak frequency and estimates it more poorly, as depicted in (b). Fig. 6 (c) examines the case when all the frequencies are well separated, and one of them is weak. The algorithm picks up a weak frequency with a smaller number of measurements when the frequencies are well separated. Taking these together, it suggests that the number of bit measurements required to accurately estimate the frequencies depends on both the frequency separations and their amplitudes. It'll be interesting to rigorously quantify this and we leave this to future work.

6.3 Performance with Finite Samples

Let $n = 100$ and $r = 3$. Assume there are a collection of T samples generated as $\boldsymbol{x}_t = \boldsymbol{U}\boldsymbol{a}_t + \boldsymbol{n}_t$, where $\boldsymbol{U}$ is an orthogonal matrix normalized from a random matrix generated with i.i.d. Gaussian entries, $\boldsymbol{a}_t \sim \mathcal{N}(0, \boldsymbol{I}_r)$ is generated with standard Gaussian entries, and $\boldsymbol{n}_t \sim \mathcal{N}(0, \sigma^2\boldsymbol{I}_n)$ is independently generated Gaussian entries. All m sensors measure the same set of T samples and communicate their

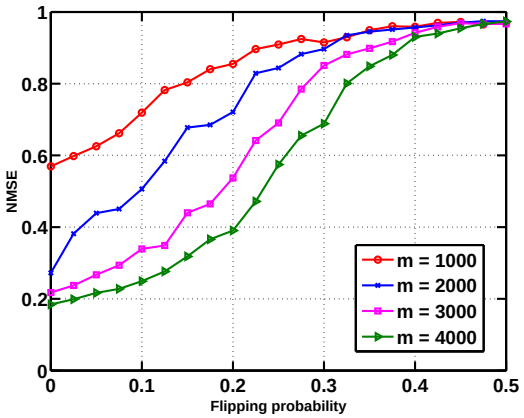

Figure 4: Performance with flipping errors: NMSE with respect to the number of bit measurements with different flipping probabilities when $n = 100$ and $r = 3$.

bit measurements for subspace estimation. Our method can be regarded as subspace estimation through comparison measurements without directly observing the original data samples, which may be useful in a crowdsourcing application.

Fig. 7 (a) shows the NMSE of principal subspace estimate with respect to the number of bit measurements for different number of samples $T = 10, 20, 30, 40, 50, 100$ and 200 averaged over 20 Monte Carlo runs when the samples are noise-free with $\sigma^2 = 0$. As T increases, the NMSE decreases as the bit measurements get more accurate in light of Theorem 2. Note that the performance gain diminishes as T is sufficiently large as all bit measurements are accurate with high probability. This changes as the data samples are contaminated by noise as in Fig. 7 (b), when $\sigma^2 = 0.1$. It is evident that more samples are necessary for the aggregation procedure to yield accurate bit measurements, and performance improves as more samples are averaged. We leave the theoretical analysis with noisy samples for future work.

7 Concluding Remarks

In this paper, we present a simple distributed sensing and central estimation framework to recover the principal subspace of a low-rank covariance matrix from a near-optimal number of one-bit energy comparisons, and described the sample complexity to ensure the one-bit measurement is accurate. We also develop an online subspace estimation algorithm to ease the memory requirement at the fusion center with a slight increase in communication overhead. Numerical examples are provided to validate the proposed approach. In the future work, we will extend the approach to enable faster tracking of the principal subspace [1], and apply to spectrum sensing in cognitive radio applications. Finally, it is of interest to extend the observation model to the sparse logistic regression framework in [20] and the more general link functions in [9].

Acknowledgement

The author thanks Lee Seversky, Lauren Huie and Matt Berger for their hospitality and helpful discussions during her stay at the Air Force Research Lab, Rome, New York where part of this work was accomplished. The author also thanks Yuxin Chen for helpful discussions regarding Lemma 1.

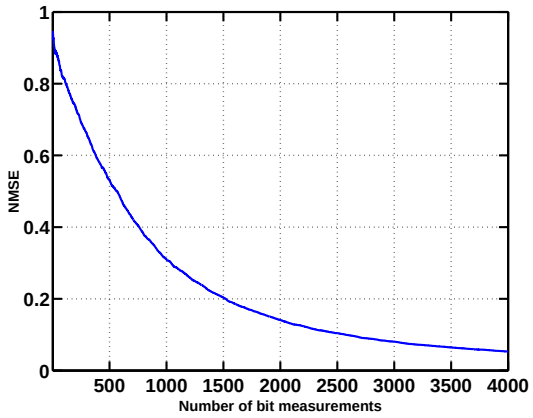

Figure 5: The NMSE of principal subspace with respect to the number of bit measurements when $n = 40$ and $r = 3$ in an online setting. The estimation accuracy is comparable to that in a batch setting.

This work was partially supported by the Air Force Summer Faculty Fellowship Program and the Google Faculty Research Award.

Appendix

A Supporting Lemmas

In the appendix, we first record several inequalities that are useful in our proof.

Lemma 3 (scalar Bernstein's inequality with sub-exponential norm, [21]). *Let $z_1, \ldots, z_L$ be independent random variables with $\mathbb{E}[z_k] = 0$ and $\sigma_k^2 = \mathbb{E}[z_k^2]$, and $\mathbb{P}[|z_k| > u] \leq Ce^{-u/\sigma_k}$ for some constants C and σ_k. Define $\sigma^2 = \sum_{k=1}^L \sigma_k^2$ and $B = \max_{1\leq k\leq L} \sigma_k$. Then*

$$\mathbb{P}\left[\left|\sum_{k=1}^L z_k\right| > u\right] \leq 2\exp\left(-\frac{u^2}{2C\sigma^2 + 2Bu}\right).$$

Lemma 4 (matrix Bernstein's inequality with sub-exponential norm, [22]). *Let $\boldsymbol{X}_1, \cdots, \boldsymbol{X}_L$ be independent zero-mean symmetric random matrices of dimension $n\times n$. Suppose $\sigma^2 = \left\|\sum_{k=1}^L \mathbb{E}[\boldsymbol{X}_k \boldsymbol{X}_k^T]\right\|$ and $\|\boldsymbol{X}_k\|_{\psi_1} \leq B$ almost surely for all k. Then for any $\tau > 0$,*

$$\mathbb{P}\left[\left\|\sum_{k=1}^L \boldsymbol{X}_k\right\| > \tau\right] \leq 2n\exp\left(-\frac{\tau^2}{2\sigma^2 + 2B\tau/3}\right). \tag{12}$$

Lemma 5 (Davis-Kahan, [15]). *Let $\boldsymbol{A}, \tilde{\boldsymbol{A}} \in \mathbb{R}^{n\times n}$ be symmetric matrices. Denote $\boldsymbol{V}$ as the subspace spanned by the top r eigenvalues of $\boldsymbol{A}$, and $\tilde{\boldsymbol{V}}$ as the subspace spanned by the top r eigenvalues of $\tilde{\boldsymbol{A}}$. Let δ be the spectral gap between the rth and the $(r+1)$th eigenvalue of $\boldsymbol{V}$. Then there exists a $r \times r$ orthogonal matrix $\boldsymbol{Q}$ such that the two subspaces $\boldsymbol{V}$ and $\tilde{\boldsymbol{V}}$ is bounded by*

$$\|\tilde{\boldsymbol{V}} - \boldsymbol{V}\boldsymbol{Q}\|_{\mathrm{F}} \leq \frac{2\sqrt{2r}\|\tilde{\boldsymbol{A}} - \boldsymbol{A}\|}{\delta}.$$

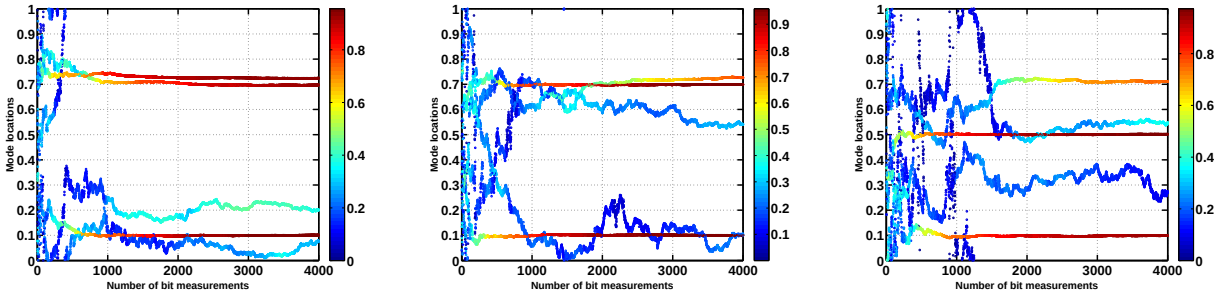

(a) $\mathcal{F} = [0.1, 0.7, 0.725]$; $\sigma^2 = [1, 1, 1]$;

(b) $\mathcal{F} = [0.1, 0.7, 0.725]$; $\sigma^2 = [1, 1, 0.5]$;

(c) $\mathcal{F} = [0.1, 0.5, 0.7]$; $\sigma^2 = [1, 1, 0.3]$;

Figure 6: Online line spectrum estimation: the estimated frequency values against the number of bit measurements when $n = 40$ and $r = 3$. Each time 5 frequencies are estimated with the color bar indicates their amplitudes. The true frequency profiles and their amplitudes are given in the subtitles.

Lemma 6 (Hanson-Wright inequality, [23]). *Let $\mathbf{\Sigma}$ be a fixed $n \times n$ matrix. Consider a random vector $\boldsymbol{x} = (X_1, \ldots, X_n)$ where X_i are independent random variables satisfying $\mathbb{E}X_i = 0$ and $\|X_i\|_{\psi_2} \leq K$. Then for any $t \geq 0$, we have*

$$\mathbb{P}\left[\left|\boldsymbol{x}^T\mathbf{\Sigma}\boldsymbol{x} - \mathbb{E}\boldsymbol{x}^T\mathbf{\Sigma}\boldsymbol{x}\right| > t\right] \leq 2\exp\left[-c\min\left(\frac{t^2}{K^4\|\mathbf{\Sigma}\|_{\mathrm{F}}^2}, \frac{t}{K^2\|\mathbf{\Sigma}\|}\right)\right].$$

B Proof of Lemma 1

Proof. We first prove (7). For $k = 1, \ldots, n - r$,

$$\begin{aligned}
\boldsymbol{v}_k^T\bar{\boldsymbol{J}}\boldsymbol{v}_k &= \mathbb{E}\left[\frac{1}{m}\sum_{i=1}^m \operatorname{sign}\left(\sum_{k=1}^r \lambda_k(\boldsymbol{a}_i^T\boldsymbol{u}_k)^2 - \sum_{k=1}^r \lambda_k(\boldsymbol{b}_i^T\boldsymbol{u}_k)^2\right)\left((\boldsymbol{a}_i^T\boldsymbol{v}_k)^2 - (\boldsymbol{b}_i^T\boldsymbol{v}_k)^2\right)\right] \\
&= \mathbb{E}\left[\operatorname{sign}\left(\sum_{k=1}^r \lambda_k(\boldsymbol{a}_i^T\boldsymbol{u}_k)^2 - \sum_{k=1}^r \lambda_k(\boldsymbol{b}_i^T\boldsymbol{u}_k)^2\right)\left((\boldsymbol{a}_i^T\boldsymbol{v}_k)^2 - (\boldsymbol{b}_i^T\boldsymbol{v}_k)^2\right)\right], \\
&= \mathbb{E}\left[\operatorname{sign}\left(\sum_{k=1}^r \lambda_k(\boldsymbol{a}_i^T\boldsymbol{u}_k)^2 - \sum_{k=1}^r \lambda_k(\boldsymbol{b}_i^T\boldsymbol{u}_k)^2\right)\right]\mathbb{E}\left[(\boldsymbol{a}_i^T\boldsymbol{v}_k)^2 - (\boldsymbol{b}_i^T\boldsymbol{v}_k)^2\right] = 0,
\end{aligned}$$

where the penultimate equation follows from that $\boldsymbol{a}_i^T\boldsymbol{v}_k$'s and $\boldsymbol{a}_i^T\boldsymbol{u}_k$'s are independent from the Gaussianity of $\boldsymbol{a}_i$, and the last equality follows from $\mathbb{E}\left[(\boldsymbol{a}_i^T\boldsymbol{v}_k)^2 - (\boldsymbol{b}_i^T\boldsymbol{v}_k)^2\right] = 0$.

We next prove (6). For $k = 1, \ldots, r$,

$$\begin{aligned}
\boldsymbol{u}_k^T\bar{\boldsymbol{J}}\boldsymbol{u}_k &= \mathbb{E}\left[\operatorname{sign}\left(\sum_{k=1}^r \lambda_k(\boldsymbol{a}_i^T\boldsymbol{u}_k)^2 - \sum_{k=1}^r \lambda_k(\boldsymbol{b}_i^T\boldsymbol{u}_k)^2\right)\left((\boldsymbol{a}_i^T\boldsymbol{u}_k)^2 - (\boldsymbol{b}_i^T\boldsymbol{u}_k)^2\right)\right] \\
&:= \mathbb{E}\left[\operatorname{sign}\left(\sum_{k=1}^r \lambda_k V_k\right)V_k\right],
\end{aligned}$$

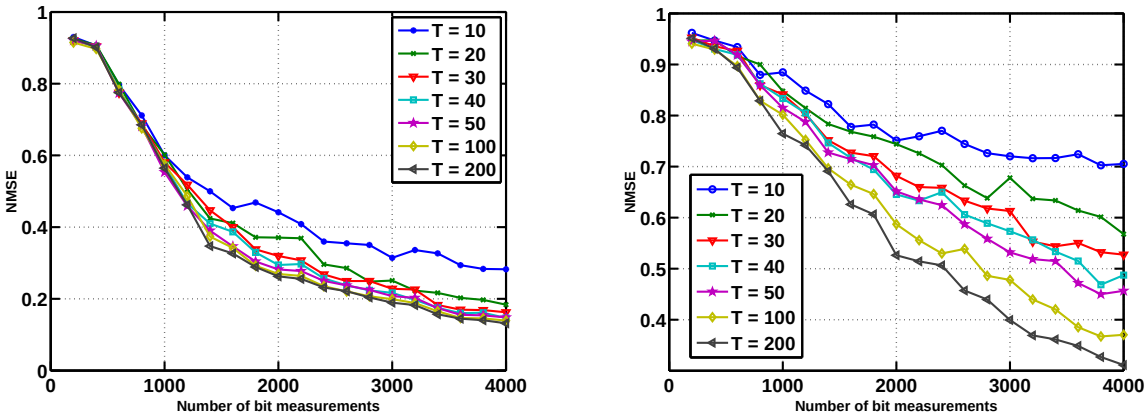

(a) $\sigma^2 = 0$ (b) $\sigma^2 = 0.1$

Figure 7: The NMSE of principal subspace estimate with respect to the number of bit measurements when $n = 40$ and $r = 3$ for different number of samples for (a) noise-free samples and (b) noisy samples.

where $V_k = (\boldsymbol{a}_i^T \boldsymbol{u}_k)^2 - (\boldsymbol{b}_i^T \boldsymbol{u}_k)^2$'s are i.i.d. random variables following the Laplace distribution with parameter 1. Let $E = \sum_{k' \neq k} \lambda_{k'} V_{k'}$, then by iteratively applying conditional expectation,

$$\begin{aligned} \boldsymbol{u}_k^T \bar{\boldsymbol{J}} \boldsymbol{u}_k &= \mathbb{E}\left[\mathbb{E}\left[\operatorname{sign}\left(\lambda_k V_k + E\right) V_k | E = \epsilon\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\operatorname{sign}\left(V_k + \epsilon/\lambda_k\right) V_k | E = \epsilon\right]\right]. \end{aligned}$$

First, assume $\epsilon > 0$,

$$\begin{aligned} \mathbb{E}\left[\operatorname{sign}\left(V_k + \epsilon/\lambda_k\right) V_k | E = \epsilon\right] &= \int_{-\infty}^{-\epsilon/\lambda_k} (-v) f_{V_k}(v) dv + \int_{-\epsilon/\lambda_k}^{\infty} v f_{V_k}(v) dv \\ &= 2 \int_{\epsilon/\lambda_k}^{\infty} v f_{V_k}(v) dv \\ &= (1 + \epsilon/\lambda_k) e^{-\epsilon/\lambda_k}. \end{aligned}$$

Similarly we derive it for $\epsilon < 0$, together we have

$$\mathbb{E}\left[\operatorname{sign}\left(V_k + \epsilon/\lambda_k\right) V_k | E = \epsilon\right] = (1 + |\epsilon|/\lambda_k) e^{-|\epsilon|/\lambda_k}.$$

It is straightforward that $(1 + |\epsilon|/\lambda_k) e^{-|\epsilon|/\lambda_k} \leq 1$, therefore $\boldsymbol{u}_k^T \bar{\boldsymbol{J}} \boldsymbol{u}_k \leq 1$. On the other hand,

$$\begin{aligned} \boldsymbol{u}_k^T \bar{\boldsymbol{J}} \boldsymbol{u}_k &= \mathbb{E}\left[(1 + |E|/\lambda_k) e^{-|E|/\lambda_k}\right] \\ &\geq \mathbb{E}[e^{-|E|/\lambda_k}] = \mathbb{E}\left[e^{-\left|\sum_{k' \neq k} \lambda_{k'} V_{k'}\right|/\lambda_k}\right]. \end{aligned} \tag{13}$$

Next, we provide two lower bounds on (13), then (6) follows by taking the maximum of the two

bounds. The first lower bound follows straightforwardly from

$$\begin{aligned}\mathbb{E}\left[e^{-|\sum_{k'\neq k}\lambda_{k'}V_{k'}|/\lambda_k}\right] &\geq \mathbb{E}\left[e^{-\sum_{k'\neq k}\lambda_{k'}|V_{k'}|/\lambda_k}\right] \\ &= \prod_{k'\neq k}\mathbb{E}\left[e^{-\lambda_{k'}|V_{k'}|/\lambda_k}\right] \\ &= \prod_{k'\neq k} 2\int_0^\infty e^{-\lambda_{k'}v/\lambda_k} f_{V_{k'}}(v)dv = \prod_{k'\neq k}\frac{\lambda_k}{\lambda_k+\lambda_{k'}}. \end{aligned} \tag{14}$$

Combining (13) and (14), we obtain

$$\boldsymbol{u}_k^T\bar{\boldsymbol{J}}\boldsymbol{u}_k \geq \prod_{k'\neq k}\frac{\lambda_k}{\lambda_k+\lambda_{k'}} = \prod_{k'\neq k}\frac{1}{1+\lambda_{k'}/\lambda_k} \geq \left(\frac{1}{1+\kappa(\boldsymbol{\Sigma})}\right)^{r-1}.$$

The second lower bound follows from

$$\begin{aligned}\mathbb{E}\left[e^{-|\sum_{k'\neq k}\lambda_{k'}V_{k'}|/\lambda_k}\right] &= \int_0^\infty \mathbb{P}\left[e^{-|\sum_{k'\neq k}\lambda_{k'}V_{k'}|/\lambda_k}\geq h\right]dh \\ &= \int_0^\infty \mathbb{P}\left[\left|\sum_{k'\neq k}\frac{\lambda_{k'}}{\lambda_k}V_{k'}\right| \leq \log\left(\frac{1}{h}\right)\right]dh \\ &\geq \mathbb{P}\left[\left|\sum_{k'\neq k}\frac{\lambda_{k'}}{\lambda_k}V_{k'}\right| \leq \kappa(\boldsymbol{\Sigma})\right]e^{-\kappa(\boldsymbol{\Sigma})}\end{aligned}$$

where the last equation is obtained by setting $h = e^{-\kappa(\boldsymbol{\Sigma})}$. Applying Lemma 3 with $u = \kappa(\boldsymbol{\Sigma})$, $B = \sqrt{2}\kappa(\boldsymbol{\Sigma})$, $C = 1$ and $\sigma^2 = 2r\kappa^2(\boldsymbol{\Sigma})$, we have

$$\begin{aligned}\mathbb{P}\left[\left|\sum_{k'\neq k}\frac{\lambda_{k'}}{\lambda_k}V_{k'}\right| \leq \kappa(\boldsymbol{\Sigma})\right] &\geq 1-2\exp\left(-\frac{\kappa^2(\boldsymbol{\Sigma})}{4r\kappa^2(\boldsymbol{\Sigma})+2\sqrt{2}\kappa(\boldsymbol{\Sigma})}\right) \\ &\geq 1-\exp\left(-\frac{1}{8r}\right) \geq \frac{1}{9r}.\end{aligned}$$

where the last inequality follows from $\exp\left(-\frac{1}{8r}\right) \leq 1-\frac{1}{8r}+\frac{1}{128r^2} \leq 1-\frac{1}{9r}$. Combined with (13), we obtain

$$\boldsymbol{u}_k^T\bar{\boldsymbol{J}}\boldsymbol{u}_k \geq \frac{e^{-\kappa(\boldsymbol{\Sigma})}}{9r}.$$

□

C Proof of Lemma 2

Proof. We write

$$\boldsymbol{J}-\bar{\boldsymbol{J}} = \frac{1}{m}\sum_{i=1}^m\left[y_i\boldsymbol{W}_i-\bar{\boldsymbol{J}}\right] := \frac{1}{m}\sum_{i=1}^m\boldsymbol{B}_i,$$

where $\boldsymbol{B}_i = y_i\boldsymbol{W}_i - \bar{\boldsymbol{J}}$. To apply Lemma 4, we have $\mathbb{E}[\boldsymbol{B}_i] = 0$, and

$$\begin{aligned}\|\boldsymbol{B}_i\| &\leq \|y_i(\boldsymbol{a}_i\boldsymbol{a}_i^T-\boldsymbol{b}_i\boldsymbol{b}_i^T)\| + \|\bar{\boldsymbol{J}}\| \\ &\leq \|\boldsymbol{a}_i\|_2^2 + \|\boldsymbol{b}_i\|_2^2 + 1\end{aligned}$$

where the last inequality follows from $\|\bar{\boldsymbol{J}}\| \leq 1$ using Lemma 1. It is obvious that both $\|\boldsymbol{a}_i\|_2^2$ and $\|\boldsymbol{b}_i\|_2^2$ are chi-squared random variables with degrees of freedom n, then $\boldsymbol{B}_i$ is sub-exponential with bounded sub-exponential norm. We have $\|\boldsymbol{B}_i\|_{\psi_1} \leq Cn$ for some constant C. For the variance, we need to compute $\mathbb{E}[\boldsymbol{B}_i^T \boldsymbol{B}_i]$:

$$\begin{aligned}\mathbb{E}[\boldsymbol{B}_i^T \boldsymbol{B}_i] &= \mathbb{E}[(y_i(\boldsymbol{a}_i\boldsymbol{a}_i^T - \boldsymbol{b}_i\boldsymbol{b}_i^T) - \bar{\boldsymbol{J}})^T (y_i(\boldsymbol{a}_i\boldsymbol{a}_i^T - \boldsymbol{b}_i\boldsymbol{b}_i^T) - \bar{\boldsymbol{J}})] \\ &= \mathbb{E}[(\boldsymbol{a}_i\boldsymbol{a}_i^T - \boldsymbol{b}_i\boldsymbol{b}_i^T)(\boldsymbol{a}_i\boldsymbol{a}_i^T - \boldsymbol{b}_i\boldsymbol{b}_i^T)] - \bar{\boldsymbol{J}}^T\bar{\boldsymbol{J}} \\ &= 2\mathbb{E}[\boldsymbol{a}_i\boldsymbol{a}_i^T\|\boldsymbol{a}_i\|_2^2] - 2\boldsymbol{I} - \bar{\boldsymbol{J}}^T\bar{\boldsymbol{J}} = n\boldsymbol{I} - \bar{\boldsymbol{J}}^T\bar{\boldsymbol{J}},\end{aligned}$$

therefore $\sigma^2 = \left\|\sum_{i=1}^m \mathbb{E}[\boldsymbol{B}_i^T\boldsymbol{B}_i]\right\| \leq \sum_{i=1}^m \left\|\mathbb{E}[\boldsymbol{B}_i^T\boldsymbol{B}_i]\right\| = m\max\{n, \|\bar{\boldsymbol{J}}^T\bar{\boldsymbol{J}}\|\} = mn$. Applying Lemma 4 we have,

$$\mathbb{P}\left[\left\|\frac{1}{m}\sum_{i=1}^m \boldsymbol{B}_i\right\| > \tau\right] \leq 2n\exp\left(-\frac{\tau^2 m}{2mn^2 + 2Cn\tau/3}\right).$$

Rearranging will conclude the proof. □

D Proof of Proposition 1

Proof. Conditioned on $\boldsymbol{a}_i$ and $\boldsymbol{b}_i$, we have $\mathbb{E}[z_i|\boldsymbol{a}_i, \boldsymbol{b}_i] = \langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle$. Let

$$z_i - \mathbb{E}[z_i|\boldsymbol{a}_i, \boldsymbol{b}_i] = \frac{1}{T}\sum_{t=1}^T \langle \boldsymbol{W}_i, \boldsymbol{\Sigma} - \boldsymbol{x}_t\boldsymbol{x}_t^T\rangle := \frac{1}{T}\sum_{t=1}^T Q_t,$$

where $Q_t = \langle \boldsymbol{W}_i, \boldsymbol{\Sigma} - \boldsymbol{x}_t\boldsymbol{x}_t^T\rangle$. We may appeal to the Bernstein-type inequality in Lemma 4. First, $\mathbb{E}[Q_t|\boldsymbol{a}_i, \boldsymbol{b}_i] = 0$. Second,

$$\begin{aligned}\mathrm{Var}[Q_t|\boldsymbol{a}_i, \boldsymbol{b}_i] &= \mathrm{Var}[\langle \boldsymbol{W}_i, \boldsymbol{x}_t\boldsymbol{x}_t^T\rangle|\boldsymbol{a}_i, \boldsymbol{b}_i] \\ &= \mathrm{Var}\left[(\boldsymbol{a}_i^T\boldsymbol{x}_t)^2 - (\boldsymbol{b}_i^T\boldsymbol{x}_t)^2|\boldsymbol{a}_i, \boldsymbol{b}_i\right] \\ &\leq 2\mathrm{Var}[(\boldsymbol{a}_i^T\boldsymbol{x}_t)^2|\boldsymbol{a}_i] + 2\mathrm{Var}[(\boldsymbol{b}_i^T\boldsymbol{x}_t)^2|\boldsymbol{b}_i] \\ &= 2(\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i + \boldsymbol{b}_i^T\boldsymbol{\Sigma}\boldsymbol{b}_i) := 2B.\end{aligned}$$

Third, since

$$\begin{aligned}|Q_t| &\leq |\langle \boldsymbol{W}_i, \boldsymbol{\Sigma} - \boldsymbol{x}_t\boldsymbol{x}_t^T\rangle| \\ &\leq \left|\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i - \boldsymbol{b}_i^T\boldsymbol{\Sigma}\boldsymbol{b}_i - (\boldsymbol{x}_t^T\boldsymbol{a}_i)^2 + (\boldsymbol{x}_t^T\boldsymbol{b}_i)^2\right| \\ &\leq \boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i + \boldsymbol{b}_i^T\boldsymbol{\Sigma}\boldsymbol{b}_i + |\boldsymbol{x}_t^T\boldsymbol{a}_i|^2 + |\boldsymbol{x}_t^T\boldsymbol{b}_i|^2,\end{aligned}$$

where $\boldsymbol{x}_t^T\boldsymbol{a}_i|\boldsymbol{a}_i \sim \mathcal{N}(0, \boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i)$ and $\boldsymbol{x}_t^T\boldsymbol{b}_i|\boldsymbol{b}_i \sim \mathcal{N}(0, \boldsymbol{b}_i^T\boldsymbol{\Sigma}\boldsymbol{b}_i)$. Since $|\boldsymbol{x}_t^T\boldsymbol{a}_i|^2/\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i$ and $|\boldsymbol{x}_t^T\boldsymbol{b}_i|^2/\boldsymbol{b}_i^T\boldsymbol{\Sigma}\boldsymbol{b}_i$ are exponential random variables with parameter 1, we have

$$\|Q_t\|_{\psi_1} \leq 2(\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i + \boldsymbol{b}_i^T\boldsymbol{\Sigma}\boldsymbol{b}_i) := 2B.$$

Assume $\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle$ is positive without loss of generality, we have

$$\mathbb{P}\left[|z_i - \mathbb{E}[z_i|\boldsymbol{a}_i, \boldsymbol{b}_i]| > |\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle||\boldsymbol{a}_i, \boldsymbol{b}_i\right] \leq 2\exp\left(-\frac{|\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle|^2 T}{4B(1 + |\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle|/3)}\right) \tag{15}$$

Next, we can bound the quadratic form $\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i$ using the Hanson-Wright inequality [23] in Lemma 6. Since $\mathbb{E}[\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i] = \mathrm{Tr}(\boldsymbol{\Sigma})$, with probability at least $1-\delta/3$, we have

$$\left|\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i - \mathrm{Tr}(\boldsymbol{\Sigma})\right| \le c\|\boldsymbol{\Sigma}\|_{\mathrm{F}}\log(1/\delta)$$

for some constant c. Since $\|\boldsymbol{\Sigma}\|_F \le \mathrm{Tr}(\boldsymbol{\Sigma})$, we have $\boldsymbol{a}_i^T\boldsymbol{\Sigma}\boldsymbol{a}_i \le c_1\mathrm{Tr}(\boldsymbol{\Sigma})\log(1/\delta)$ with probability at least $1-\delta/3$ for some absolute constant c_1. Denote this as event $\mathcal{G}_1$.

Further from the arguments in [3, Proposition 1], we have that

$$c_2\|\boldsymbol{\Sigma}\|_{\mathrm{F}}\log(1/\delta) \le |\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle| \le c_3\|\boldsymbol{\Sigma}\|_{\mathrm{F}}\log(1/\delta)$$

with probability at least $1-\delta/3$ for some absolute constants c_2 and c_3. Denote this as event $\mathcal{G}_2$.

Conditioned on the event $\mathcal{G}_1$ and $\mathcal{G}_2$, and plug in the above into (15), we have as soon as

$$T \ge c_4\frac{\mathrm{Tr}(\boldsymbol{\Sigma})}{\|\boldsymbol{\Sigma}\|_{\mathrm{F}}}\log^2(1/\delta) \tag{16}$$

for some constant c_4, the RHS of (15) can be upper bounded by $\delta/3$. To summarize, assuming (16) holds,

$$\begin{aligned}
&\mathbb{P}\left[z_i \ne \mathrm{sign}(\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle)\right] \\
&\le \int \mathbb{P}\left[z_i \ne \mathrm{sign}(\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle)|\boldsymbol{a}_i, \boldsymbol{b}_i\right] d\mu(\boldsymbol{a}_i)d\mu(\boldsymbol{b}_i) \\
&\le \mathbb{P}(\mathcal{G}_1^c) + \mathbb{P}(\mathcal{G}_2^c) + \int_{\mathcal{G}_1,\mathcal{G}_2} \mathbb{P}\left[|z_i - \mathbb{E}[z_i|\boldsymbol{a}_i, \boldsymbol{b}_i]| > |\langle \boldsymbol{W}_i, \boldsymbol{\Sigma}\rangle||\boldsymbol{a}_i, \boldsymbol{b}_i\right] d\mu(\boldsymbol{a}_i)d\mu(\boldsymbol{b}_i) \le \delta.
\end{aligned}$$

Our proposition then follows. □

References

[1] Y. Chi, Y. C. Eldar, and R. Calderbank, "Petrels: Parallel subspace estimation and tracking by recursive least squares from partial observations," *Signal Processing, IEEE Transactions on*, vol. 61, no. 23, pp. 5947–5959, 2013.

[2] L. Balzano, R. Nowak, and B. Recht, "Online identification and tracking of subspaces from highly incomplete information," *Proc. Allerton 2010*, 2010.

[3] Y. Chen, Y. Chi, and A. Goldsmith, "Exact and stable covariance estimation from quadratic sampling via convex programming," October 2013. [Online]. Available: http://arxiv.org/abs/1310.0807

[4] Y. Chen, Y. Chi, and A. J. Goldsmith, "Estimation of simultaneously structured covariance matrices from quadratic measurements," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, May 2014, pp. 7669–7673.

[5] ——, "Robust and universal covariance estimation from quadratic measurements via convex programming," in *IEEE International Symposium on Information Theory*, July 2014.

[6] P. T. Boufounos and R. G. Baraniuk, "1-bit compressive sensing," in *Information Sciences and Systems, 2008. CISS 2008. 42nd Annual Conference on.* IEEE, 2008, pp. 16–21.

[7] L. Jacques, J. Laska, P. Boufounos, and R. Baraniuk, "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *Information Theory, IEEE Transactions on*, vol. 59, no. 4, pp. 2082–2102, April 2013.

[8] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1275–1297, 2013.

[9] ——, "Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach," *Information Theory, IEEE Transactions on*, vol. 59, no. 1, pp. 482–494, 2013.

[10] D. Romero and G. Leus, "Wideband spectrum sensing from compressed measurements using spectral prior information," *Signal Processing, IEEE Transactions on*, vol. 61, no. 24, pp. 6232–6246, Dec 2013.

[11] D. D. Arianandа and G. Leus, "Compressive wideband power spectrum estimation," *Signal Processing, IEEE Transactions on*, vol. 60, no. 9, pp. 4775–4789, 2012.

[12] O. Mehanna and N. Sidiropoulos, "Frugal sensing: Wideband power spectrum sensing from few bits," *Signal Processing, IEEE Transactions on*, vol. 61, no. 10, pp. 2693–2703, May 2013.

[13] ——, "Adaptive thresholding for distributed power spectrum sensing," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, May 2013, pp. 4459–4463.

[14] Y. Mroueh and L. Rosasco, "Quantization and greed are good: One bit phase retrieval, robustness and greedy refinements," *arXiv preprint arXiv:1312.1830*, 2013.

[15] J. Lei and A. Rinaldo, "Consistency of spectral clustering in sparse stochastic block models," *arXiv preprint arXiv:1312.2050*, 2013.

[16] M. Brand, "Incremental singular value decomposition of uncertain data with missing values," in *ECCV 2002*. Springer, 2002, pp. 707–720.

[17] ——, "Fast low-rank modifications of the thin singular value decomposition," *Linear Algebra and its Applications*, pp. 20–30, 2006.

[18] U. Grenander and G. Szegő, *Toeplitz forms and their applications*. Univ of California Press, 1958.

[19] R. Roy and T. Kailath, "ESPRIT-estimation of signal parameters via rotational invariance techniques," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 37, no. 7, pp. 984 –995, Jul 1989.

[20] M. A. Davenport, Y. Plan, E. v. d. Berg, and M. Wootters, "1-bit matrix completion," *arXiv preprint arXiv:1209.3672*, 2012.

[21] A. W. Van Der Vaart and J. A. Wellner, *Weak Convergence*. Springer, 1996.

[22] J. A. Tropp, "User-friendly tail bounds for sums of random matrices," *Foundations of Computational Mathematics*, vol. 12, no. 4, pp. 389–434, 2012.

[23] M. Rudelson and R. Vershynin, "Hanson-Wright inequality and sub-Gaussian concentration," *arXiv preprint arXiv:1306.2872*, June 2013.